

February 13, 2008

Re: WC Docket No. 07-52

To: Marlene H. Dortch

Office of the Secretary

Federal Communications Commission

445 12th Street, SW, Suite TW-A325

Washington, DC 20554.

From: George Ou

235 2nd St, San Francisco, CA 94105

Dear Chairman Martin and Commissioners,

This is a comment on the two items in WC Docket 07-52, the Petition for Declaratory Ruling by Free Press, et. al., and the Petition for Rulemaking by Vuze. Inc.

I come to you as a concerned private citizen, a former network engineer who designed TCP/IP networks and server infrastructure, an Editor of a technology publication who has written extensively¹ on these matters and was cited by Comcast in their response to the FCC on these matters. My views do not necessarily reflect the views of my employer ZDNet (CNET Networks) since I do not represent my employer.

Background

Independent groups last year found that Comcast was sending TCP RESET packets to BitTorrent seeders at various times of the day to cut back the number of upload sessions they could have. A BitTorrent seeder is someone who is not downloading but acting as a dedicated and peer-to-peer file server. BitTorrent downloads or uploads while downloading were not affected. Various groups complained that this was possibly illegal protocol discrimination using forged TCP RESET packets while Comcast maintained that this was reasonable network management to assure fair distribution of bandwidth to all their users.

The upstream contention problem

A typical Cable broadband network such as Comcast operates under the DOCSIS 1.1 standard which offers 10 mbps of upstream bandwidth and 40 Mbps of downstream bandwidth shared amongst the neighborhood. Since the typical user has a static upstream cap of 384 kbps, it would be possible for 26 BitTorrent seeders and/or BitTorrent uploaders to completely jam the upstream pipe rendering the entire network unbearable. Since a typical Cable broadband company provisions between 50 and 400 users (typically somewhere in the middle) per cable loop, it is possible for ~10% of the users can jam the entire upstream network which ultimately affects downloads as well since services can't be asked for. This is further complicated by the fact that DOCSIS networks use a reservation system for upstream traffic on a collision network. Too many requests for upload slots and the requests collide and no one gets to transmit anything.

Accusations of discrimination

Some have complained that this was content discrimination. But Comcast does not discriminate based on content; Comcast discriminates against excessive upstream usage that chokes up their entire broadband network. The EFF complains that this was "protocol discrimination" against BitTorrent and other P2P (peer-to-peer) applications, but it is a fact that BitTorrent and P2P are the biggest upstream bandwidth users. Since BitTorrent seeders who only continuously upload throughout the day can be reasonably classified as dedicated servers, they actually fall under prohibited services under Comcast's TOS (Terms Of Service) ².

Blocking versus delaying

Comcast says they're merely delaying BitTorrent seeders from uploading to their peers while their critics say they are blocking. It is true that Comcast blocks BitTorrent seeds when the broadband network is very busy, but they do allow BitTorrent seeding at most other times of the day. Network Engineer and Internet pioneer Richard Bennett explained this best in his comment to the FCC that since BitTorrent and other peer-to-peer applications all have the ability to resume transmission at where they left off, temporary blocking of seeders effectively acts as a delaying mechanism. The file eventually gets served to the remote party outside of Comcast's network intact.

Consumer versus commercial Internet connection

The reality is that Comcast customers were never blocked, throttled, or delayed from receiving any services; they were delayed from offering hosting services (BitTorrent seeding) that were technically prohibited to begin with under the terms of service. Comcast's consumer broadband service technically doesn't have to act as a commercial hosting service to other customers in and outside of Comcast's network so the fact that they permit seeding most of the day seems like a reasonable compromise. Furthermore, BitTorrent users who are downloading are continuously uploading during the download without any delaying action so it isn't as if Comcast refuses to participate in P2P uploads.

Blocking of Lotus Notes

Comcast's network management mechanisms did have a bug in them that accidentally blocked Lotus Notes traffic, but this issue was fixed months ago when the issue was first brought to the attention of Comcast. All software and hardware implementations have bugs and we expect the service provider to act in good faith and repair the problems as soon as possible. In this particular case, Comcast appears to have acted quickly and properly by fixing the problems that blocked Lotus Notes.

The complaint to the FCC

The Free Press and Public Knowledge filed a formal complaint to the FCC to immediately enjoin Comcast from these network management practices before the merits are decided and the facts weighed. This is an unreasonable request since Comcast customers would be harmed by network traffic jams due to the lack of any traffic management. The Free Press and Public Knowledge also demanded fines of \$195,000 per infraction which would amount to over \$2

trillion dollars if we counted every Comcast customer. This is obviously impossible since it exceeds the gross revenue of any corporation in the USA.

Impractical alternatives proposed

The Free Press and Public Knowledge acknowledged the need for Comcast to manage their network and they offered some alternatives. Here is a quote from their petition to the FCC³:

"More importantly, if Comcast is concerned that the collective set of users running P2P applications are affecting quality of service for other users on a cable loop, they could readily set dynamic quotas for each user on the loop, so as to ensure that there is always bandwidth available for users who are not running P2P applications – and they could do so without interfering in protocol choice. Or they could also charge by usage, provide more bandwidth to all users, or actually offer high symmetric broadband speeds."

The first proposal from the Free Press and Public Knowledge petition is the use of "dynamic quotas for each user on the loop". This would certainly be an ideal solution, the only problem is that such a mechanism does not exist. Since we are dealing with upload congestion, the place to solve the problem would be the cable modems at the customer premise. The problem can't deal with this at the head unit CMTS (Cable Modem Termination System) since the damage would already be done to the shared Cable medium before the flood of upstream data ever reached the CMTS. The currently deployed cable modems only have static upstream quotas that can be configured upon boot up, but they cannot be changed dynamically on the fly.

The second proposal from the Free Press and Public Knowledge petition is to "charge for usage" or in other words, a metered Internet service. The problem is that the Free Press is hypocritical on this proposal since they are actively criticizing Time Warner's metered Internet service. In this Free Press press release⁴, they state:

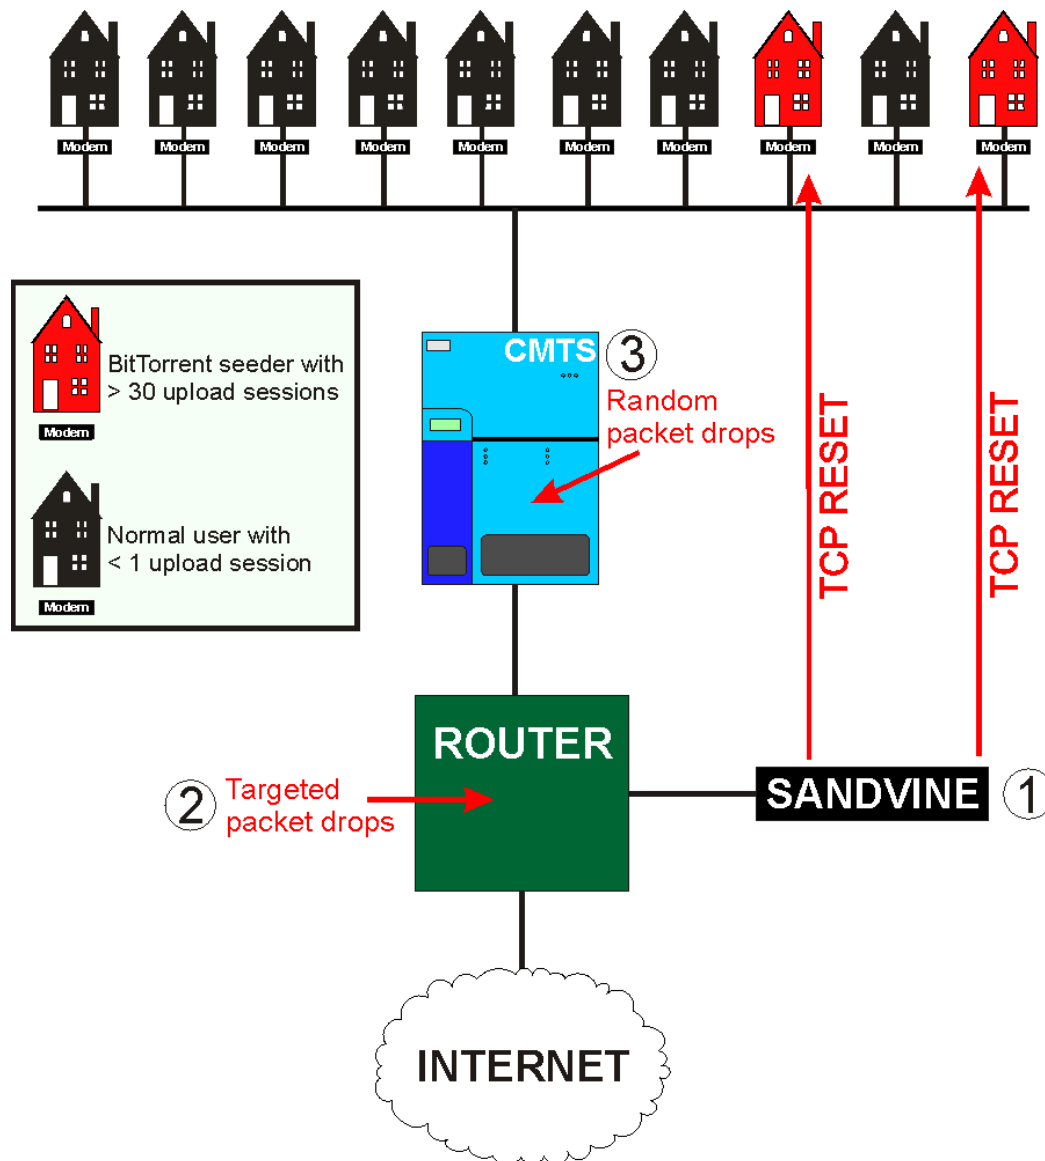
"Compared to that approach, Time Warner's proposal is better — at least customers will know what they're getting into. But metered prices may chill innovation in cutting-edge applications because consumers will have a disincentive to use them. Viewed in the context of our long-term national goals for a world-class broadband infrastructure, telling consumers they must choose between blocking and metered pricing is a worrying development."

While I'm in full agreement that metered Internet service is a bad idea⁵, I am shocked that Free Press, Public Knowledge, and even the EFF⁶ would propose metered Internet service which is something that is so anti-consumer. It's even more shocking how inconsistent and hypocritical the Free Press is. Note that the EFF did not join the Free Press and Public Knowledge FCC petition, but they echoed many of the same things in their paper "Packet Forgery By ISPs: A Report on the Comcast Affair"⁷. So the very services such as video distribution over BitTorrent that the Free Press, Public Knowledge, and the EFF are asking to be protected are the very services that would be completely shut down under a metered Internet service model. It would be simply be cost prohibitive to run any sort of P2P applications.

The final proposal from the Free Press and Public Knowledge petition is that Comcast should "actually offer high symmetric broadband speeds". Now everyone wants faster Internet service, but would Comcast's inability to deliver immediately on this constitute a crime and deserve a fine? If anything, a large fine would prevent them from reaching their publicly stated goal of 20% conversion to DOCSIS 3.0 which would offer a 12-fold increase in upload and a 4-fold increase in download speeds (120 Mbps up and 160 Mbps down). So what's actually driving this conversion to DOCSIS 3.0? Verizon's FiOS and AT&T's U-Verse service putting competitive pressure on Comcast in a free market, not Free Press and Public Knowledge.

Finding the fair and practical network management solution

The EFF and others have suggested that packet dropping would be a better way to solve the congestion problems and it avoids using forged TCP RESET packets but this is neither practical nor fair. To illustrate this, I've come up with the following illustration with three technical solutions to the problem of upstream congestion with varying efficacy.



The illustration above is a simplified architectural view of a DOCSIS shared-medium network. Note that the packet switched portion of the network we call the Internet doesn't really start until the data gets to the CMTS. That means traditional packet switching traffic engineering doesn't apply on the first hop since the reservation system works on a collision network. Each house in the drawing represents approximately 10 to 20 homes. The red homes represent the heavy BitTorrent users who seed and act as dedicated BitTorrent servers with tens of upstream

connections while the black homes represent the typical household which uses relatively few upstream connections.

Solution 1: Use TCP RESET packets to reduce the excessive upload sessions from the relatively few users acting as BitTorrent seeds. Even though we can't say for certain that Sandvine is the exact appliance being used to send the TCP resets, it doesn't matter since this is the general method that Comcast uses. Since BitTorrent is engineered to grab as much bandwidth in either direction as it can with explicit design goals to thwart network management mechanisms, no formal application-level throttling mechanism is available. That means the only mechanism that can be employed to throttle BitTorrent is via external informal methods and the TCP RESET is one of them. Since the application in this case doesn't understand nor would it intend to honor requests to slow down and reduce the number of upstream connections, an alternative method below the application layer is employed.

Solution 2: Use targeted packet drops to only drop or block the upstream BitTorrent packets from BitTorrent seeders. This method can work nearly as effectively as solution 1, but it's extremely expensive. A real-time ACL (Access Control List) with thousands of new throttle or block requests per second would put a massive load on any router. This would require very expensive "fork-lift upgrades" per cable loop. Solution 1 by comparison is very cheap to deploy since the Sandvine-type appliance only needs to run on commodity server and it doesn't need the cooperation of other devices to enforce the rules. Sending out TCP RESETS is fundamentally cheaper by an order of magnitude compared to targeted packet drops. It would unnecessarily increase the cost of running the network which ultimately has to be passed on to the consumer with no additional benefits whereas the money would be far better spent on the DOCSIS 3.0 upgrades.

Solution 3: Use random packet drops to reduce overall network congestion. While this method is cheap to implement and it would ultimately work by slowing every session down equally, it would be extremely unfair to the vast majority of Cable broadband subscribers. As British Telecom Chief Researcher Bob Briscoe put it in his paper to the IETF⁸ (Internet Engineering Task Force):

"It's equivalent to claiming food rations are fair because the boxes are all the same size, irrespective of how many boxes each person gets or how often they get them."

Proponents of random packet dropping are essentially saying that everyone must slowdown on upload traffic regardless of how many sessions or how little bandwidth they use by some fixed percentage. To enforce this, everyone's packets would have a certain percentage dropped which would force their TCP/IP stack to slow down. This would have disastrous consequences on protocols like VoIP since dropped packets mean dropped words within a conversation. It matters not that VoIP aren't constantly active or that they only use one session at a time and only requires a small amount of bandwidth while in use, they'll be punished just the same along with the excessive bandwidth and session users. This is a reckless approach with no sense of fairness and should be disregarded.

Conclusion

Given the fact that the petitioners don't actually dispute the right of Comcast or any other network provider to reasonably manage their network and throttle down excessive uploaders, the only thing in dispute is the methodology used to achieve these goals. I've highlighted all the proposed alternatives along with the pros and cons of each solution and I'm confident that the one already in use by Comcast is the least intrusive, least expensive, and most practical way of dealing with the realities of a DOCSIS 1.1 broadband network.

Market forces in the form of competition are forcing Comcast to upgrade to a much higher performance and more symmetrical DOCSIS 3.0 network and that is a win for the consumer. But regardless of how fast the network is, there will always be a need for practical network management solutions that ensure per-user fairness and not per-session fairness. The last thing we should do is force Comcast to implement more expensive and/or less fair traffic management schemes that at best wastes money and at worst degrade performance for consumers who are using far less than their fair share of bandwidth.

References

- 1) <http://blogs.zdnet.com/Ou/?p=852>
A rational debate on Comcast traffic management
- 2) <http://www6.comcast.net/terms/use/#prohibited>
Comcast Acceptable Use Policy For High-Speed Internet Services
- 3) http://www.freepress.net/docs/fp_et_al_nn_declaratory_ruling.pdf
In the Matter of the Petition of Free Press et al
- 4) <http://www.freepress.net/press/release.php?id=328>
Time Warner Metering Exposes America's Bigger Broadband Problems
- 5) <http://blogs.zdnet.com/Ou/?p=919>
Why metered Internet is a really bad idea
- 6) <http://blogs.zdnet.com/Ou/?p=914>
EFF wants to saddle you with metered Internet service
- 7) <http://www.eff.org/wp/packet-forgery-isps-report-comcast-affair>
Packet Forgery By ISPs: A Report on the Comcast Affair
- 8) http://www.cs.ucl.ac.uk/staff/B.Briscoe/projects/2020comms/refb/fair_ccr.pdf
Flow Rate Fairness: Dismantling a Religion